

Numerical modelling of convective processes in oceanography using adaptive meshing techniques

Flora MacTavish

November 29, 2010

Abstract

The objective is to model the restratification phase of open ocean deep convection (OODC) using the Imperial College Ocean Model (ICOM), with particular emphasis on the baroclinic vortices formed during the process. In particular a process study of a baroclinic vortex interacting with a wall or another vortex will be carried out. ICOM has an adaptive mesh so this should be better than previous methods at resolving the features of the process. Open Ocean Deep Convection is important to climate change because it helps to drive the Meridional Overturning Circulation (MOC). Any changes in OODC due that may occur due to global warming may also affect the MOC, so an improved understanding of this process is important in predicting changes in climate.

In the first section of this report salt fingers are modelled using ICOM. Salt fingers are a form of double diffusion, and they are studied here in order to validate ICOM against the work of [Özgökmen and Esenkov \(1998\)](#). Salt fingers are modelled using a structured mesh and the growth rate and flux ratio are calculated. These show close agreement with the previous numerical results. Salt fingers are then modelled on an unstructured mesh and a comparison between fixed and adaptive meshes is carried out, showing that the adaptive mesh reproduces the same results as the fixed mesh with fewer nodes. The final section concerns open ocean deep convection. The process is described and some of the techniques for modelling it are discussed. Future work plans are presented.

Contents

0.1	Parts that have been changed	1
1	Introduction	2
2	Salt fingers	4
2.1	Overview	4
2.2	The numerical model	5
2.2.1	The governing equations	5
2.2.2	The Imperial College Ocean Model	6
2.2.3	The experimental setup	8
2.3	Experimental and theoretical results	8
2.4	Results	9
2.4.1	Comparison between ICOM and Özgökmen and Esenkov (1998)	9
2.4.2	Horizontal average of salinity and temperature	10
2.4.3	Flux Ratio	11
2.5	A comparison between fixed and adaptive unstructured meshes using ICOM	13
2.5.1	A comparison between different resolution fixed meshes	14
2.5.2	Running ICOM with an adaptive mesh	14
2.5.3	Comparing fixed and adaptive runs starting from the same mesh.	15
2.5.4	Salt fingers with different wavelengths	19
2.5.5	Growing fingers with different wavelengths in ICOM	19
2.6	Discussion of results	21
3	The restratification of a water column following open ocean deep convection	23
3.1	Introduction	23
3.2	The process of Open Ocean Deep Convection	24
3.2.1	The phases of open ocean deep convection	24
3.2.2	The Labrador Sea	24
3.3	Modelling the restratification of a convection chimney	25
3.3.1	The Three Models	25
3.3.2	Experimental set-up	26
3.4	Future Plans	27
4	Appendix	30
4.1	Setting up the mesh	30
4.2	Non-dimensional parameters	30

0.1 Parts that have been changed

1, 2.2, 2.4, 3.4, 2.6.

Chapter 1

Introduction

The objective is to model the restratification phase of open ocean deep convection (OODC) using the Imperial College Ocean Model (ICOM). ICOM has an adaptive mesh, so it should be better than previous methods at resolving the features of the process. In particular we will focus on the baroclinic vortices that form, carrying out a detailed study of their interactions with a wall or each other. Open Ocean Deep Convection is important to climate change because it helps to drive the Meridional Overturning Circulation (MOC). Any changes in OODC due that may occur due to global warming may also affect the MOC, so an improved understanding of this process is important in predicting changes in climate.

The first part of this report concerns the modelling of salt fingers using ICOM. The purpose of this work is to verify ICOM against other numerical results. Salt fingering is an example of double diffusive convection; finger shapes form because heat and salt diffuse at different rates. Salt fingers are found in the ocean in the form of thermohaline staircases.

The first set of salt finger simulations have a mesh with rectangular elements and periodic boundary conditions. They are designed to reproduce previous numerical results from Özgökmen and Esenkov (1998). The fingers are stable with no secondary instability. It is shown that if a perturbation is added to the initial condition so that the fingers are slightly different lengths at the start of the simulation the instability grows and the difference in the finger lengths increases as time goes on. The flux ratio, a measure of the ratio of the potential energy gained by the temperature field to potential energy lost by the salinity field, is also calculated and agrees with Özgökmen and Esenkov (1998).

Salt fingers are then modelled on an unstructured triangular mesh. The fingers therefore have different lengths to start with, and again the instability grows throughout the simulation. These results are then compared to a run on an adaptive mesh. In adaptive runs the sizes of the elements in the mesh change throughout the simulation allowing more resolution in places where more is needed. The adaptive mesh is shown to reproduce the same results as the fixed mesh while using fewer nodes, which represents a computational saving.

Open ocean deep convection is discussed, with emphasis on the restratification phase in which the convection cell breaks down and the temperature stratification returns to a smoother stratification decreasing with depth. This phase will be modelled, starting with idealised cases. The future plans are then proposed. An idealised model on OODC restratification will be run in order to investigate the nature of the vortices that form due to the baroclinic instability. A process study of a vortex interacting with a wall will then be carried out. Mixing occurs on all scales in the ocean. Currents are responsible for large scale motions of water, wind forcing is important in the upper layers of the ocean, and tidal forcing is mixes water in areas close to the shore. Internal waves and open ocean deep convection (deep upwelling of water) are smaller scale processes that are also important in the ocean. Very small scale processes such as double diffusion, mixing due to differences in diffusivity of two density-altering components in the water, also occur, but they are not important on oceanic scales. Generally, there is stronger horizontal than vertical mixing in the ocean since vertical mixing is diapycnal; it has to occur across surfaces of constant density. This involves moving against the buoyancy force so it uses up more energy (Stewart, 2005). This report will focus on vertical mixing processes.

The evidence that vertical diapycnal mixing does occur in the ocean, despite the fact that it is more energy intensive, comes from the thermocline, the region of water near the surface that is stratified in temperature. The temperature stratification has been observed to remain constant over decades

(Munk, 1966). We would expect the bottom of the thermocline to move downwards through time due to turbulent mixing of heat in the downward direction. The fact that the thermocline remains constant suggests that there must also be upward heat transport. This upward mixing occurs mainly due to internal waves and the shear of ocean currents when they flow past particular types of topography including continental slopes, mid-ocean ridges and sea mounts. These obstacles cause the currents to become turbulent (Stewart, 2005). The eddy diffusivity is a measure of turbulent mixing. It has been measured in tracer release experiments, and the horizontal eddy diffusivity in the ocean was found to be 10^8 to 10^{10} times larger than vertical eddy diffusivity (Ledwell and Watson, 1993).

This report will discuss two types of vertical mixing. Salt fingers occur when there are opposing gradients of heat and salt, which leads to instability due to the faster diffusivity of heat (Özgökmen and Esenkov, 1998). They are small scale, and although they occur naturally in the ocean, they are not important on oceanic scales because they do not transport large volumes of water or large amounts of heat. Observations show that turbulence is more important than double diffusion to vertical mixing in the ocean (Gregg, 1991).

The other type of mixing to be discussed in this report is open ocean deep convection (OODC). This occurs at some specific convection sites, mainly in the North Atlantic. Deep water is moved up to the surface and mixed vigorously, forming convection chimneys. It helps to drive the Meridional Overturning Circulation (MOC) so it is important on oceanic scales. It is unclear how a change in the OODC process would affect the MOC. The second chapter of this report will focus on the final stage of open ocean deep convection, when the convection column breaks down and the water is restratified into its initial configuration with denser water underlying less dense water (Rousset et al., 2009).

Chapter 2

Salt fingers

2.1 Overview

Salt fingers are formed due to double diffusive convection, which is a type of mixing that occurs when water contains two dissolved components with differences in diffusivity and density. There must also be a gradient in the two components. Salt fingers are formed when a region of warmer, saltier water overlies a region of colder fresher water. When a random perturbation occurs at the boundary, the instability grows and the two regions form finger-like shapes at the boundary. This happens because heat diffuses more quickly than salt, so if a small amount of warm salty water is displaced down into the cold fresh region it will cool before it freshens, so it becomes denser, and sinks even further. This process continues and forms salt fingers. It can also happen in reverse when cold, salty water is displaced upwards it warms but its salinity remains almost the same, so it is lighter than the surroundings and continues to move upwards to produce an ascending salt finger. The scale of salt fingers is determined by a balance between thermal diffusion and viscous drag. Thermal diffusion is greatest at smaller scales and releases potential energy in the salt stratification. Viscous drag limits the smallest scales of motion, and makes the fingers smaller. Salt fingers are first discussed by [Stern \(1960\)](#), although double diffusion was already a recognised phenomenon. The salt fountain was first suggested by [Stommel \(1953\)](#). There are some regions of the ocean where there are few nutrients in the upper layers so there is little life in these regions. The salt fountain uses double diffusion to move deep, nutrient rich water to the upper regions of the ocean where photosynthesis is possible so that plankton can grow and fish can be produced as a result.

Salt fingers occur naturally in the ocean, particularly in the tropics where evaporation exceeds precipitation so the surface water is relatively warm and salty. They also occur as a result of currents meeting at particular points, bringing together water masses with different properties from different regions of the ocean. Salt fingers play a role in the formation of thermohaline staircases ([Morell and Merryfield, 2006](#)), which are vertical structures with a step-like variation in temperature and salinity (S and T). There are regions of constant S and T with thin boundary layers in between where the S and T gradients are steep. When T and S decrease with depth the staircase is characterised by salt fingering, and if T and S increase with depth it is associated with diffusive convection, another type of double diffusive process. Diffusive convection typically occurs in polar regions because the melting sea ice produces a layer of relatively cold, fresh water at the top of the ocean.

A thermohaline staircase in the western tropical North Atlantic, known as the C-SALT staircase, has been observed extensively. The warm, saline Subtropical Underwater overlies the colder, fresher, Antarctic Intermediate Water and Atlantic Central Water, creating the necessary conditions for salt fingering. [Schmitt et al. \(2005\)](#) performed a tracer release experiment to determine the structure of the staircase. There are around ten layers of uniform temperature and salinity, each between 10 and 30 m thick. In between these layers are interfaces, 0.5 – 5 m in depth, with high temperature and salinity gradients. This staircase is supported by centimetre scale salt finger convection cells at the interfaces which release the potential energy stored up by the salinity differences. They produce an unstable buoyancy flux which drives large-scale convective motions in the uniform layers, maintaining the stratification.

Salt fingers have also been implicated in the formation of an eddy in the North Atlantic ([Dmitrenko](#)

et al., 2008). An eddy with two layers was investigated, the upper layer being warmer and saltier than the lower one. the stratification is believed to be maintained by salt fingering at the interface. It was generated in the North Kara Sea from two different water masses: the Atlantic water from the Arctic Ocean, which passes through the Fram Strait, and the Atlantic water inflow from the Barents and North Kara Seas. The stratification is maintained thousands of kilometres downstream of this location. Numerical simulations suggested that motion in the upper layer is dominated by double diffusion and the lower layer is dominated by shear instability.

The results of former experimental, theoretical and numerical studies on salt fingers will be summarised in section 2.3. Section describes the numerical technique used in this report. The governing equations are presented, the Imperial College Ocean model is discussed and the experimental setup is described. The main results are presented and discussed in section 2.4. A comparison is made between the fixed and adaptive meshes and the flux ratio is calculated. Section 2.6 summarises the main conclusions from this work.

2.2 The numerical model

This section discusses the governing equations for salt fingering and how they are solved. The governing equations are introduced and non-dimensionalised in section 2.2.1 and the Imperial College Ocean Model is discussed in section 2.2.2.

2.2.1 The governing equations

The Boussinesq equations are solved by ICOM. The governing equations are the heat, salt and velocity equations, the continuity equation for incompressible flow and the linear equation of state:

$$\frac{\partial T}{\partial t} + (\mathbf{u} \cdot \nabla)T = \kappa_T \nabla^2 T \quad (2.1)$$

$$\frac{\partial S}{\partial t} + (\mathbf{u} \cdot \nabla)S = \kappa_S \nabla^2 S \quad (2.2)$$

$$\frac{\partial \mathbf{u}}{\partial t} + (\mathbf{u} \cdot \nabla)\mathbf{u} = -\nabla p - \frac{\rho}{\rho_0} g \hat{\mathbf{z}} + \nu \nabla^2 \mathbf{u} \quad (2.3)$$

$$\nabla \cdot \mathbf{u} = 0 \quad (2.4)$$

$$\rho = \rho_0 [1 - \alpha(T - T_0) + \beta(S - S_0)], \quad (2.5)$$

where \mathbf{u} is the velocity vector, ν is the kinematic viscosity, ρ and ρ_0 are the density and reference density, p is the pressure, T and T_0 are the temperature and reference temperature, S and S_0 are the salinity and reference salinity, κ_T and κ_S are the diffusivities of temperature and salt respectively, α and β are the coefficients of thermal expansion and saline contraction respectively and g is the acceleration due to gravity.

These equations are solved in two dimensions in the x-z plane, with x as the horizontal co-ordinate and z as the vertical co-ordinate. ICOM solves the above dimensional form of the equations but Özgökmen and Esenkov (1998) solves the non-dimensionalised versions. In order to obtain the non-dimensionalised form, the following substitutions are made:

$$(x, z) = L(x^*, z^*), \quad \mathbf{u} = \frac{\nu}{L^2} \mathbf{u}^*, \quad t = \frac{L}{\nu} t^*, \quad T - T_0 = \Delta T T^*, \quad S - S_0 = \Delta S S^*, \quad p - gz = \frac{\nu^2}{L^2} p^*,$$

where L is the height of the box. Here the non-starred values e.g. T denote dimensional values and the starred values e.g., T^* denote the non-dimensionalised numbers. The non-dimensionalised form of the equations is given below. Making the above substitutions and dropping stars, the following non-dimensionalised equations are obtained:

$$\frac{\partial T}{\partial t} + (\mathbf{u} \cdot \nabla)T = \frac{1}{Pr} \nabla^2 T \quad (2.6)$$

2.2. THE NUMERICAL MODEL

Table 2.1: A summary of the parameters used in the simulation. The parameters are defined in the Appendix.

Parameter	Symbol	Value
Vertical height of the domain	L	2
Horizontal length of the domain	L_x	1 or 4
Prandtl number	Pr	7
Schmitt number	Sc	700
Density ratio	R_ρ	1.8
Thermal Rayleigh number	R_T	10^6
Coefficient of thermal expansion	α	142857.0
Coefficient of saline contraction	β	79365.0
Diffusivity of temperature	κ_T	0.1429
Diffusivity of salinity	κ_S	0.001429
Acceleration due to gravity	g	1
Kinematic viscosity	ν	1
Saline Rayleigh number	R_S	4×10^6

$$\frac{\partial S}{\partial t} + (\mathbf{u} \cdot \nabla) S = \frac{1}{Sc} \nabla^2 S \quad (2.7)$$

$$\frac{\nu^2}{L^3} \frac{\partial \mathbf{u}}{\partial t} + \frac{\nu^2}{L^3} (\mathbf{u} \cdot \nabla) \mathbf{u} = -\frac{\nu^2}{L^3} \nabla p + (\alpha \Delta T T - \beta \Delta S S) g \hat{\mathbf{z}} + \frac{\nu^2}{L^3} \nabla^2 \mathbf{u} \quad (2.8)$$

In equation 2.8 the first term in the density has been absorbed by the pressure. This equation can be rewritten as

$$\frac{\partial \mathbf{u}}{\partial t} + (\mathbf{u} \cdot \nabla) \mathbf{u} = -\nabla p + \frac{R_T}{Pr} \left(T - \frac{S}{R_\rho} \right) \hat{\mathbf{z}} + \nabla^2 \mathbf{u} \quad (2.9)$$

where Pr is the Prandtl number, Sc is the Schmitt number, R_T is the thermal Rayleigh number and R_ρ is the density ratio, all defined in the Appendix.

The curl can be taken of equation 2.9 to give the equation in terms of vorticity. In two dimensions this gives

$$\frac{\partial \omega}{\partial t} + (\mathbf{u} \cdot \nabla) \omega = \frac{R_T}{Pr} \left(\frac{\partial T}{\partial x} - \frac{1}{R_\rho} \frac{\partial S}{\partial x} \right) \hat{\mathbf{y}} + \nabla^2 \omega. \quad (2.10)$$

This can be rewritten in the same form as Özgökmen and Esenkov (1998):

$$\frac{\partial \omega}{\partial t} + J(\psi, \omega) = \frac{R_T}{Pr} \left(\frac{\partial T}{\partial x} - \frac{1}{R_\rho} \frac{\partial S}{\partial x} \right) \hat{\mathbf{y}} + \nabla^2 \omega, \quad (2.11)$$

where $J(\psi, \omega) = \frac{\partial \psi}{\partial x} \frac{\partial \omega}{\partial z} - \frac{\partial \psi}{\partial z} \frac{\partial \omega}{\partial x}$. The non-dimensionalisation is used to obtain values of α , β , κ_T and κ_S so that the results will reproduce those of Özgökmen and Esenkov (1998).

2.2.2 The Imperial College Ocean Model

In this section the Imperial College Ocean model is discussed in more detail. Section 2.2.2 discusses the finite element method and section 2.2.2 discusses the adaptive mesh.

The Finite Element Method

ICOM solves the incompressible non-hydrostatic Boussinesq equations in dimensional form as shown above using the finite element method. In the finite element method the equations are written in the weak form which is obtained by multiplying the equation with a test function and integrating over the domain. An example of an equation to be solved is

$$L(u) = q, \quad (2.12)$$

2.2. THE NUMERICAL MODEL

where $L(u)$ is a linear operator. If we have an approximate solution to the equation, u^δ , then $L(u)$ does not exactly equal q , so we introduce the residual $R(u)$, which gives a measure of the error. The equation now becomes

$$R(u^\delta) = L(u^\delta) - q. \quad (2.13)$$

The next step is to multiply the equation by a test function $w(x)$ to obtain the weak form of the equation, giving

$$\int_{\Omega} w(x) R(u^\delta(x)) dx = \int_{\Omega} w(x) L(u^\delta(x)) dx - \int_{\Omega} w(x) q(x) dx \quad (2.14)$$

If $u^\delta = u$ then $R(u^\delta) = 0$. Therefore the term containing $R(u^\delta)$ is set to zero. The equation now becomes

$$\int_{\Omega} w(x) L(u^\delta(x)) dx = \int_{\Omega} w(x) q(x) dx. \quad (2.15)$$

Different projection methods use different weight functions. The solution can be represented in terms of the basis functions $N_i(x)$:

$$u(x) = \sum_{i=1}^N N_i(x). \quad (2.16)$$

In the Galerkin method the weight functions are the same set of functions as the trial functions, sp $W_j(x) = N_j(x)$.

The velocity and pressure were calculated using the continuous Galerkin discretisation (Ford et al., 2004). Continuous Galerkin methods allow only trial and test functions that are continuous between elements. Temperature was solved using the control volumes finite element discretisation. The solver was the iterative gmres method. Iterative methods solve the equations by a series of approximations that converge to the actual solution. Gmres is a krylov subspace iterative method.

Adaptivity

The mesh can be adapted as the simulation runs. The elements change size and shape so that regions of interest can be resolved in more detail. The adaptivity uses both changes in the connectivity of the mesh (h-method) and node movement (r-method). The mesh can be adapted to any field specified by the user. The mesh will be made finer in regions where the Hessian of that field is higher. This means that the mesh will be more resolved in regions where there is greater variation in the field. In the case of salt fingers the mesh should be more resolved in the fingering region, particularly at the edges of the fingers. A metric is constructed from the Hessian of the field being adapted to. The metric is then used to find the distance between the nodes in the new mesh. (Ford et al., 2004; Pain et al., 2001).

The first step is to calculate the metric. In order to do this the interpolation error must be found. In the following we assume a P^1 approximation. There is one unknown variable η . The P^1 interpolate of the solution is given by $\Pi_h \eta$. The interpolation error is therefore defined as

$$\epsilon_\infty = |\eta - \Pi_h \eta|_\infty, \quad (2.17)$$

where $|\cdot|_\infty$ denotes the L^∞ norm. The error can be related to the Hessian H of the variable η by

$$\epsilon = ||\eta - \Pi_h \eta||_{\infty, K} \leq c \cdot \max_e e^T |H| e, \quad (2.18)$$

where K is an element e is a vector ranging over element K and c is a constant. The Hessian of η is defined as

$$H = \begin{pmatrix} \frac{\partial^2 \eta}{\partial x_1^2} & \frac{\partial^2 \eta}{\partial x_1 \partial x_2} \\ \frac{\partial^2 \eta}{\partial x_1 \partial x_2} & \frac{\partial^2 \eta}{\partial x_2^2} \end{pmatrix}.$$

In the above equation $|H|$ is the absolute Hessian. The error can now written as

$$\epsilon \approx c \cdot \max_e e^T |H| e. \quad (2.19)$$

If ϵ is the target error in the L_∞ norm, equation 2.19 can be written as

$$1 \approx \max_e e^T \left(\frac{c}{\epsilon} |H| \right) e, \quad (2.20)$$




Figure 2.1: The initial conditions for the temperature and salinity fields. They are both 1 in the top half of the domain and 0 in the bottom half. The temperature field has a region in the middle of height 0.2 with a linear gradient. The salt field is similar, but there is also a sinusoidal perturbation at the interface given by $0.016 \{\cos(10\pi x)\}$. This perturbation forces the salt fingers to form. These initial conditions are designed to reproduce the results of Özgökmen (Özgökmen and Esenkov, 1998).

The metric used in producing the mesh is therefore $M = \frac{c_0}{\epsilon} H$, where ϵ is the interpolation error specified by the user. The mesh satisfying the metric is the one that gives a unit edge length for all the edges with respect to the metric M .

2.2.3 The experimental setup

Salt fingers were modelled in two dimensions using the same parameters as Özgökmen and Esenkov (1998) in order to compare their results to ICOM. The simulation was run on a mesh with horizontal and vertical dimensions given by $L_x = 4$ by $L_z = 2$, where x is the horizontal coordinate-ordinate and y is the vertical coordinate-ordinate. Both temperature and salinity were initially set to 0 in the bottom half of the domain and 1 in the top half, with an intermediate zone of height 0.2 where the fields had a linear gradient (figure 2.1). An initial perturbation in temperature or salinity is necessary for salt finger formation; this was created by adding $A \{\cos(10\pi x)\}$ to the initial salinity configuration (figure 2.1), where A is the amplitude of the initial perturbation. The heat and salt were not allowed to flow through the boundaries.

There are two slightly different set ups. In section ?? the model is set up in exactly the same way as Özgökmen and Esenkov (1998), with a 1 by 2 mesh with square elements and $A = 0.016$. The sides have periodic boundaries and the top and bottom have Dirichlet boundaries. In section 2.5 the domain is 4 by 2 and an unstructured triangular grid is used and $A = 0.02665$. This is to test the adaptivity in ICOM. All sides have Dirichlet boundary conditions in this case.

2.3 Experimental and theoretical results

Some of the numerical studies of salt fingers and their main results are summarised in this section. Salt fingers were first modelled numerically in 1980 by Piasek and Toomre (Piasek, 1980). This work and subsequent developments in the field are summarised in Yoshida and Nagashima (2003). Piasek and Toomre (1980) used two-dimensional $x - z$ grids. They produced plots of the salinity field as the salt fingers evolved, which clearly showed the shape of eight salt fingers, including the blobs on the ends of the fingers that had been observed in the laboratory. The faster diffusion of heat was clear because the heat profiles were smoother.

Whitfield et al. (1989) used a Prandtl number between 3 and 10 and a Lewis number between 0.1 and 0.5 (both defined in section 4.2). This does not correspond to heat-salt or sugar-salt fingers because the ratios of the saline to thermal diffusivity and thermal diffusivity to kinematic viscosity do not correspond to the ratios for salt-heat or sugar-salt solutions. The values for salt-heat experiments

are a Prandtl number of 7 and a Lewis number of $\frac{1}{100}$. They carried out five different simulations with the density ratio between 1.1 and 2.5. They resolved fingers and the blobs on the ends, but their flux ratio was different from the theoretical and experimental values. The flux ratio is a measure of the ratio of potential energy gained by the temperature field to the potential energy lost by the salinity field. This was attributed to the large Stern number which meant that shear instability did not apply. The Stern number is the ratio of the net vertical heat flux and the viscous dissipation. When it is greater than $\frac{1}{3}$ the salt fingers become unstable (Holyer, 1981).

Shen has carried out several simulations of salt fingers (Shen and Veronis, 1997). The evolution of salt fingers corresponded more closely to laboratory experiments in the later stages of the simulation, with salt fingers and convective plumes both present, and the interface remaining in the same place. Unlike Whitfield et al, Shen reproduced the correct flux ratio. The formation of large-scale eddies in the upper and lower sections was also demonstrated.

Özgökmen et al (Özgökmen and Esenkov, 1998; Yoshida and Nagashima, 2003) used a two-dimensional rectangular box with width $L_x = 1$ and height $L_z = 2$. The Prandtl number was set to 7 and the Schmitt number was 700 (both defined in section 4.2). The density ratio was 1.8 and the Rayleigh number, the ratio of buoyancy and viscosity forces, was 1×10^{-6} . The resolution of the box was 201 grid points in each direction. The ratio of thermal to saline diffusivity $\frac{K_T}{K_S}$ was given by the ratio of the Schmitt number to the Prandtl number $\frac{Sc}{Pr} = 100$ and the density ratio R_ρ was 1.8 (defined in section 4.2).

Salt fingers were modelled for both linear and nonlinear equations of state. In the nonlinear version there is an extra nonlinear term in temperature. This takes into account the non-linear effects due to the temperature, which causes the interface between the layers to move upwards in laboratory experiments. The nonlinearity is characterised by the parameter ϵ

The simulation was run for ϵ values of 0 (which corresponds to a linear equation of state), 1 (which has some nonlinearity) and 3 (larger nonlinear term). Plots were produced of the upper and lower interfaces between the fingering zone and the mixed region against time. The linear case became unstable before the nonlinear case. The lower interface moved more rapidly than the upper one in the nonlinear simulation but the upper and lower interfaces moved at the same rate in the linear one. The finger zone thickness and growth rate were plotted against time for each value of ϵ . The finger zone thickness increased linearly and then exponentially. The thickness increased more quickly for the linear case. The growth rate initially follows the theoretical model (Schmitt, 1979) but then departs from it. Schmitt's work was nonlinear so only the ϵ case could be applied. This deviation is explained by the fact that Schmitt omits the z -dependence and assumes constant background temperature and salinity gradients. The vertical kinetic energy was plotted against time for ascending and descending fingers for each value of ϵ . For finite ϵ the kinetic energy of the descending fingers was higher than that of the ascending fingers. For zero ϵ the ascending and descending fingers had the same kinetic energy curves.

2.4 Results

The results are presented in this section. First, in section 2.4.1, a comparison is presented between ICOM and Özgökmen and Esenkov (1998). The purpose of this is to see whether ICOM can reproduce these results from another model. In section salt fingers are modelled using ICOM with different resolutions. In section adaptive results are presented and compared to fixed mesh results. The purpose of this is to see whether ICOM reproduces the same results with an adaptive mesh as with a fixed mesh. If so, the adaptivity is a good way to save computing power without compromising on the results.

2.4.1 Comparison between ICOM and Özgökmen and Esenkov (1998)

The results from ICOM were compared to the results obtained by Özgökmen and Esenkov (1998). This was done by running ICOM using the same type of mesh as that used in Özgökmen and Esenkov (1998). A fixed, structured mesh with rectangular elements, of 100 by 200 elements was used. The domain was 1 by 2.

Figure 2.2 shows the temperature and salinity fields at time 0.05. The salinity field displays the characteristic structure of the salt fingers, with 'blob' shapes on the ends that form due to self-organisation of the flow. The salinity field is more diffused. The fingers are all the same length, with no secondary instability. This is due to the fact that the mesh is structured.




Figure 2.2: The salinity (left) and temperature (right) fields at time=0.05. The fingers are all the same length: there is no secondary instability because a structured mesh is used. The salinity field displays that characteristic shape of salt fingers with "plumes" developing at the top of the fingers. The temperature field is more diffused.




Figure 2.3: The y-coordinate of the top of the upward growing fingers and the bottom of the downward growing fingers is plotted. The ends of the fingers are taken to be the interface between the finger region and the mixed region. The results from ICOM are plotted along with the results of Özgökmen and Esenkov (1998) for comparison. The ICOM results show good agreement with Özgökmen and Esenkov (1998).

The interfaces between the salt fingers and the mixed zone are plotted as a function of time in figure 2.3. The ICOM results are plotted with the results of Özgökmen and Esenkov (1998) for comparison. There is good agreement between the results. Both the ICOM results and the Özgökmen and Esenkov (1998) results exhibit stable growth. The fingers start off the same length as each other and remain the same length as each other until the end of the simulation. This is because a structured mesh was used. This means that the fingers started off exactly the same length so there was no opportunity for secondary instability to grow. In section 2.5 an unstructured mesh is used and a secondary instability grows because of the initial differences in the finger lengths.

2.4.2 Horizontal average of salinity and temperature

Figure 2.4 shows the horizontal averages of temperature and salinity fields for a fixed mesh run. The temperature/salinity fields were split into layers of height Δy , and all the temperature/salinity values were averaged within each section. The result is plotted as a function of y . This graph can be compared to figure ?? which shows the temperature and salinity fields at this time. Both curves go from 0 at the bottom of the domain to 1 at the top. The salinity curve reflects the shape of the salt fingers. The salinity is 1 at the top of the domain. It then reduces down to a minimum. This is the blob region, where there is more fresh water. The region just below the blob has more salty water, so the salinity




Figure 2.4: The average temperature and salinity as a function of vertical co-ordinate. The average is calculated by splitting the domain into strips of the same width as the domain and height dz where dz is small compared to the height of the domain. The temperature and salinity values were averaged within each section. The temperature field is highly diffused but the salinity field reflects the shape of the salt fingers, with a minimum at the top of the upward growing fingers and a maximum at the bottom of the downward growing fingers.

goes back up to a maximum. It then gradually decreases along the finger length until it reaches another minimum. This is the region just above the blob, where there is more fresh water. The blob region has more salty water, so there is another maxima. The salinity then goes down to zero below the fingering region. The temperature curve is a smooth curve because temperature is more diffusive than salinity. The temperature goes from 1 at the top to 0 at the bottom. The change is fairly smooth but it is steeper at the top and bottom of the fingers. At the top of the fingers the average is over the top of the fluid which has a temperature of 1 and the top of the fingers which have a temperature of 0, so there is a sudden drop in temperature as at the beginning of the finger zone.

2.4.3 Flux Ratio

The Flux ratio is the ratio of potential energy gained by the temperature field to the potential energy lost by the salinity field. The first step to calculating the flux ratio is to express temperature, salinity and vertical velocity as

$$w = \bar{w} + w' \quad (2.21)$$


$$T = \bar{T} + T' \quad (2.22)$$

$$S = \bar{S} + S' \quad (2.23)$$


where \bar{w} is the horizontal mean of the vertical velocity component (ie the mean velocity at a given height y), and w' is the deviation from the horizontal mean of the vertical component of the velocity at a particular point. Similarly, \bar{T} is the horizontal mean of the temperature, and T' is the deviation from the mean; and \bar{S} is the horizontal mean of the salinity field and S' is the deviation from the mean.

The deviations from the mean are presented in figure 2.5. They are calculated as a function of y . The domain was split into sections of height Δy and the mean temperature (salinity) calculated in each section as above. The standard deviation from the mean was calculated for each section. The deviations from the mean in the salinity field (figure 2.5(b)) reaches a maxima in both blob regions for both upward and downward growing fingers. Figure 2.2(a) shows that the blob regions have a higher degree of variability than the surrounding regions. The double vortex structure in the plumes increases the variability and leads to overturning. There is also a maximum just above the middle of the fingers for upward growing fingers and just below the middle for downward growing fingers. The figure also shows greater variability here.

The temperature variability also has maxima just above (below) the middle for upward (downward) growing fingers. This is confirmed by figure 2.2(b) which shows the most variability. The temperature has a smooth increase in variability from the outside of the fingering region to the middle of the region.



(a) Deviations from the mean for the temperature field as a function of y coordinate-ordinate.



(b) Deviations from the mean for the salinity field as a function of y coordinate-ordinate.

Figure 2.5: Deviations from the horizontal mean for temperature and salinity at time 0.04s. The temperature deviations reach a maximum in the middle of the domain where the temperature shows a lot of variation. The salinity deviations reach their maximum values in the blob regions at the ends of the fingers.

This is again confirmed by figure 2.2(b) and makes sense because heat is more diffusive than salt. Because there is no mean vertical velocity, $w' = w$ (Özgökmen and Esenkov, 1998). This is plotted in figure 2.6. w was plotted as a function of y by averaging all the w values in each horizontal section. The velocity is highest in the blob regions because the space in between the blobs is small.

The advective heat and salt fluxes can now be expressed as

$$F_T = \overline{w'T'} \quad F_S = \overline{w'S'}. \quad (2.24)$$

The advective heat (salt) flux is calculated for a given range of y values by finding the deviation from the horizontal mean for the temperature (salinity) at each x value. Each of these is multiplied with the deviation from the horizontal mean of the vertical velocity at that point. These sums are then averaged to obtain the heat (salt) flux for that y value. This is done for all the horizontal sections to find the heat (salt) fluxes as a function of y . The flux ratio is given by

$$R_F = \frac{\alpha F_T}{\beta F_S}. \quad (2.25)$$

This can now easily be calculated from F_T and F_S . In order to calculate the diagnostics for an adaptive mesh, the salinity, temperature and velocity fields are transposed onto a structured mesh with the same resolution at each point. The diagnostics can then be calculated in the same way.




Figure 2.6: The average velocity field for upward and downward fingers. The velocity reaches a maximum at the ends of the fingers.




Figure 2.7: The flux ratio from ICOM and Özgökmen and Esenkov (1998). These show good agreement.

As the flux ratio is the ratio of potential energy gained by the temperature field to the potential energy lost by the salinity field it can be used to deduce the level of fingering activity. R_F must be below unity, otherwise the system would gain energy overall. The diffusive processes transport more heat than salt, and advective processes transport more salt than heat. The fingers themselves are largely formed by the advection of salt. Therefore the flux ratio gives the ratio of diffusive processes to advective processes. It approaches unity when advective transport is low and diffusive transport is high. A low flux ratio denotes more advective transport and therefore more finger activity. Flux ratios in laboratory experiments vary from 0.4 to 0.8 (Özgökmen and Esenkov, 1998). The flux ratio here is in that range.

The flux ratio has only been plotted for the fingering region. High values were obtained for the flux ratio outside the fingering region. The flux ratio should be zero in these regions, but zero was not obtained because finite values were obtained for the heat and salt fluxes in these regions, possibly due to numerical errors. The flux ratio was high because $F_T \approx 10^{-13}$ and $F_S \approx 10^{-24}$, so $F_T \ll F_S$. This may be due to the heat being more diffusive than the salt. This seems unphysical, the flux ratio should be zero outside the fingering region because there is little transport of heat or salt.

2.5 A comparison between fixed and adaptive unstructured meshes using ICOM

In this section the results obtained by the Imperial College Ocean Model (ICOM) will be analysed. Adaptive and fixed meshes are compared. The aim is to show that the adaptive mesh can reproduce the

same results as the fixed mesh with fewer nodes. Several diagnostics are also calculated to analyse the properties of the salt fingers. Three different resolution fixed meshes are compared in section 2.5.1. It is shown that there is a secondary instability. The magnitude of the secondary instability can be linked to the resolution of the mesh, and therefore can be used to compare between fixed and adaptive meshes. Fixed and adaptive runs with different initial meshes are compared in section 2.5.2. Fixed and adaptive meshes with the same initial mesh are compared in section 2.5.3; this allows a more rigorous comparison between the fixed and adaptive meshes. Different initial conditions are compared in section 2.5.4 and section 2.5.5 contains results with a linear gradient in temperature and salinity across the whole domain rather than a sharp interface.

2.5.1 A comparison between different resolution fixed meshes

Three different resolution fixed mesh simulations were run for the same initial condition, a wavelength of 0.2. The meshes had roughly 150000 nodes, 300000 nodes and 600000 nodes, and all had unstructured triangular elements. The three runs are shown in figure ???. Plume shaped dipolar vortices are seen forming on the ends of the fingers; these play a role in limiting the length of the salt fingers because they eventually come off the end (Taylor and Veronis, 1996). These dipolar vortices are a common feature of two-dimensional fluids and are believed to form due to the self-organisation of the flow (Özgökmen and Esenkov, 1998). These plumes make the fingers wider at the ends so that interaction occurs between them. This causes a secondary instability which can be seen in figure ??? from 0.04s. At this time the fingers are different lengths from their neighbours, forming a periodic instability. This instability appears to be larger in the lower resolution runs.

The secondary instability was measured by finding the standard deviation of the lengths of the salt fingers from the mean at several time steps. The length of the upward and downward salt fingers was calculated by splitting the domain into sections the width of each salt finger. The length of the salt finger was taken to be the contour where salinity is 0.5. The top of the salt finger was therefore the maximum y coordinate-ordinate where the salinity was equal to 0.5 for each section. The lengths of all the salt fingers were averaged for each time step and the standard deviation was calculated. All the calculations are made over the middle 75% of the domain in order to reduce the edge effects unless otherwise stated.

The finger lengths provide a validation to show that the different resolutions are producing the same results. The lengths of the fingers as a function of time are plotted in figure 2.8(b). This shows that the fingers grow at the same rate for all the resolutions. Figure 2.8(a) shows that the 600,000 node run has the lowest standard deviation of the three and the 150,000 node one has the highest standard deviation. This shows that the instability grows fastest in the lower resolution runs. Part of the reason for this is that the initial condition is less resolved in the lower resolution runs. This means that there is initially a larger difference in the finger heights for lower resolution runs. This instability grows exponentially so the lower resolution runs have a larger instability at later times. The instability is also partly caused by the Dirichlet boundary conditions at the sides of the domain. There is more instability at the edges than the middle. However, this effect was largely overcome by calculating the diagnostics over the middle of the domain only.

2.5.2 Running ICOM with an adaptive mesh

Salt fingers were also modelled using adaptive meshes to change the mesh so that particular areas are more finely resolved than others. In this instance the initial mesh was a structured mesh. This allowed the initial condition to be better resolved as the fingers all started with exactly the same length.

The mesh was adapted to the salinity field, which means that the nodes were made smaller in the areas where the curl of the salinity was greater. The area around the fingers, and particularly the interface between the fingers and the rest of the mesh are given a higher resolution. It is possible to adapt to other parameters such as velocity. In this case it is not worth adapting to other parameters because most of the resolution is needed in the fingering area, and this can be achieved just by adapting to salinity. This was checked by adapting to velocity as well and then checking whether there was any significant difference in the meshes produced. Adapting to velocity did not appear to make any difference to the result.

A comparison between three different adaptive runs is shown in figures 2.9(a) and 2.9(b). The three adaptive runs have different interpolation error bounds for the salinity field. This the the largest error

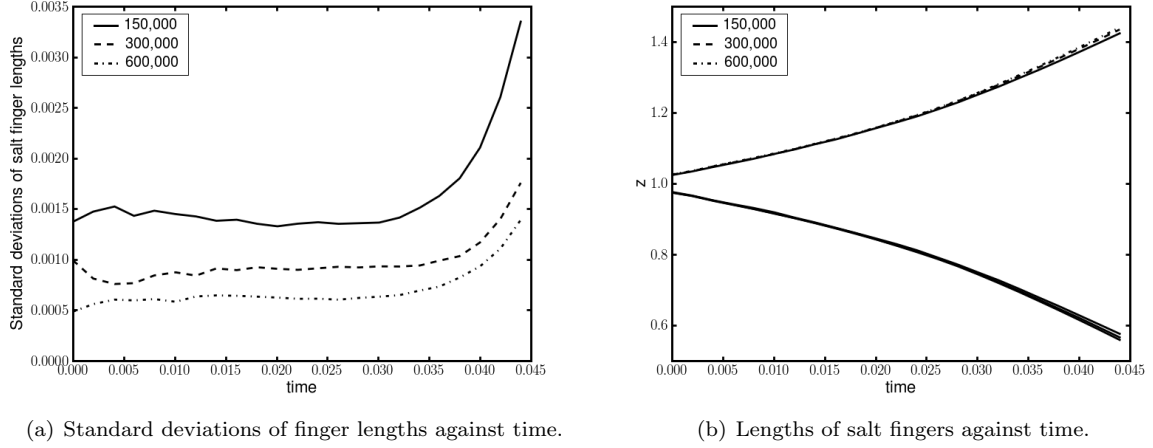


Figure 2.8: Results for different resolution fixed meshes. The three runs produce the same results in terms of the average finger lengths, but there is a difference in the standard deviation of the finger lengths. The standard deviation of the finger lengths is larger for lower resolution runs because the secondary instability grows more quickly in the lower resolution runs.

that is allowed in the salinity field when the salinity is being interpolated from the old mesh onto the new mesh. The lower the number, the higher the resolution of the mesh will be in areas of high salinity gradient. This can be seen from the number of nodes: the run with interpolation error of 0.04 has the fewest nodes and the one with 0.01 has the most nodes. The graph of the standard deviations shows that the lowest resolution run clearly has the lowest standard deviation, as expected. This is consistent with the idea that the lower the resolution, the higher the standard deviation. The other two are less conclusive. This could be because of the way that the initial condition is resolved.

Figure 2.9(d) compares an adaptive run with two fixed runs. In both cases the adaptive run performs better than the fixed runs, despite having fewer nodes. Some of this advantage is due to a better resolution of the initial condition in the adaptive runs. This is partly because the initial condition is applied after the mesh is adapted for the first time, so there is a lot more resolution in that region when the initial condition is applied. Also the adaptive run starts from a structured mesh. When the mesh adapts for the first time it maintains some of the structure and the fingers are almost exactly the same length initially. The initial deviation from the mean is therefore smaller than in the fixed runs, which are unstructured meshes with the same resolution everywhere. The initial standard deviation of the finger lengths is therefore lower in the adaptive runs. The deviations grow exponentially, so if the adaptive mesh has smaller deviations at the beginning it will probably have smaller deviations throughout the simulation.


Although this is not a direct comparison because of the difference in applying the initial conditions, the adaptive mesh technique produces a simulation with a higher effective resolution than a fixed mesh run with the same number of nodes. Both these adaptive runs perform better than the fixed run with 600,000 nodes, so a mesh with more than 600,000 nodes would be needed to reproduce the results, although the adaptive run used less than 100,000 nodes by the end of the simulation.

2.5.3 Comparing fixed and adaptive runs starting from the same mesh.

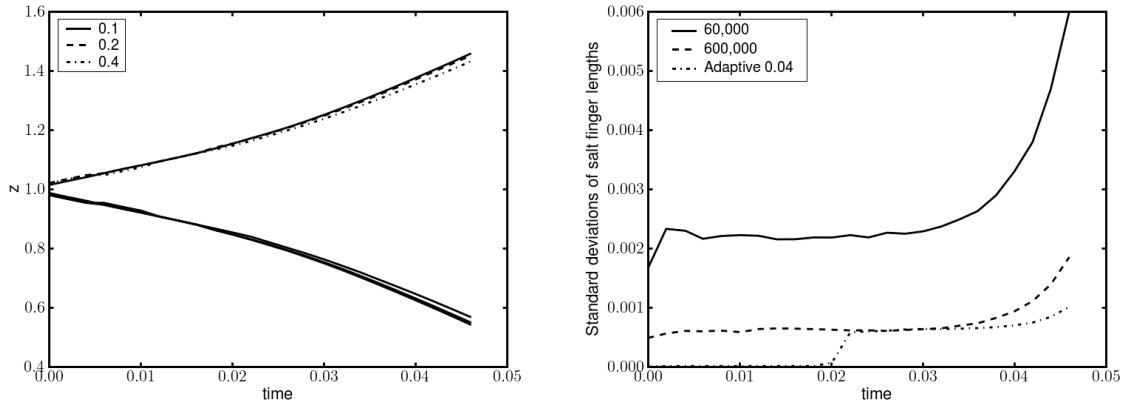
In order to make a direct comparison between a fixed and adaptive run, an adaptive simulation was run starting from the 600,000 node mesh used for the fixed mesh run in section 2.5.2. The mesh was not adapted before the first time step, so the initial condition was resolved in the same way for the fixed and adaptive runs, allowing a better comparison between the two. The interpolation error was varied as a function of time. Initially it was set to 0.00001; it then increases linearly.

Figure 2.11 shows a comparison between the fixed run with 600,000 nodes and the adaptive run which starts from this mesh. The very low interpolation error at the beginning meant that it did not lose too much resolution in the first part of the run. The adaptive run had fewer nodes throughout the

2.5. A COMPARISON BETWEEN FIXED AND ADAPTIVE UNSTRUCTURED MESHES USING ICOM



(a) The number of nodes against time for the three adaptive runs. The numbers given are the interpolation error bounds; a lower interpolation error bound gives a higher resolution mesh. (b) Standard deviations of the finger lengths for the three adaptive runs. The two lower resolution runs perform particularly well because they use fewer nodes than the lowest resolution adaptive run but have lower standard deviations.



(c) Positions of the tops and bottoms of the fingers for one adaptive run with interpolation error 0.4 and several fixed runs. This provides a validation to show that the adaptive runs get the same results as the fixed runs. (d) Standard deviations of the finger lengths for the 0.04 adaptive run and two fixed runs. This is the lowest resolution adaptive run, but it performs better than either of the fixed runs. This is partly because it uses a structured mesh to start with, which means that the initial condition gets resolved better than in the fixed runs.

Figure 2.9: A comparison between three adaptive runs with interpolation errors of 0.04, 0.02 and 0.01 (labelled 0.04, 0.02 and 0.01 respectively), and some fixed runs labelled with the number of nodes. The adaptive runs have a lower standard deviation of the finger lengths, implying that the secondary instability grows more quickly. However, the adaptive runs start off from a structured mesh and the fixed runs are unstructured so the initial condition is better resolved in the adaptive runs.




Figure 2.10: The salinity field for different resolutions (fixed mesh). Number of nodes increases left to right: 150,000; 300,000 and 600,000 nodes. Time increases top to bottom: 0.03, 0.04 and 0.05. There is a secondary instability that forms as time progresses: the fingers gradually become different lengths from their neighbours. This is due to the interactions between the fingers which occurs partially due to the plume shaped structures on the ends of the fingers. The lower the resolution, the faster the instability forms. This is used to compare fixed and adaptive meshes.

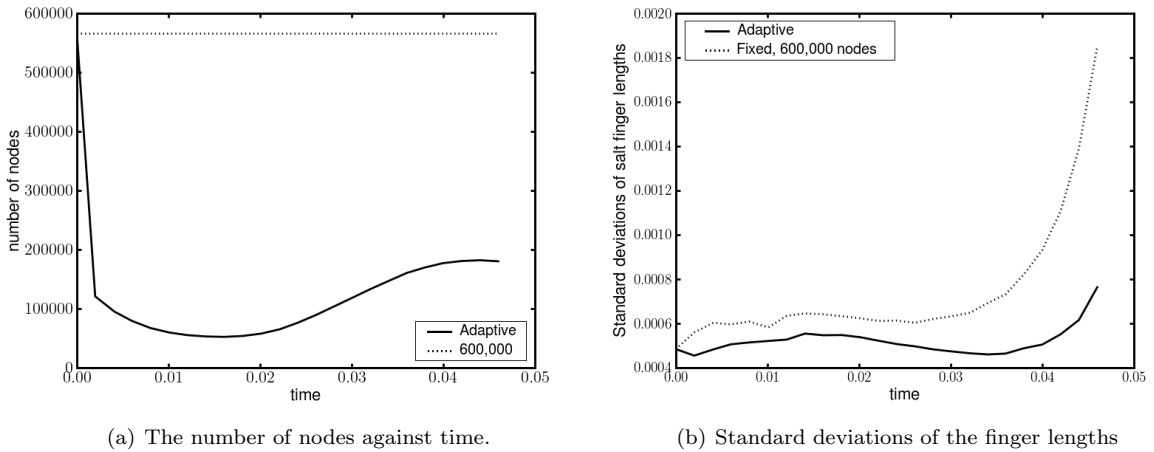


Figure 2.11: Comparison between a fixed and adaptive run, both starting from the same mesh, with no adapt at the first time step. The adaptive run has lower standard deviations, showing that the secondary instability grows more slowly in the adaptive run, so the effective resolution of the adaptive run is higher than the fixed run. Since the initial condition is resolved in the same way for both runs, this result shows that the adaptive run is definitely performing better than the fixed run.

run, and also had a lower standard deviation. This shows that the adaptive run performs better than a fixed run with fewer nodes even when the initial condition is resolved in exactly the same way.

2.5.4 Salt fingers with different wavelengths

In the previous sections the fingers were forced to grow with a particular wavelength. This was chosen to reproduce previous work and because it is close to the fastest growing mode. Theoretically there is a fastest growing mode which would emerge naturally if the initial condition was random. A formula has been obtained for the fastest growing mode by [Kunze \(1987\)](#). The wavelength λ of the fastest growing mode is given by

$$\lambda = 2\pi L \left[\frac{L}{h} R_T \left(1 - \frac{1}{R_\rho} \right) \right]^{-1/4}, \quad (2.26)$$

where h is the height of the fingering interface. The length scale L is the same length scale that appears in the Rayleigh number, and was introduced when the formula above was rewritten in terms of the Rayleigh number. The relevant length scale here is the height of the domain.

This formula was obtained for the case where there was a linear temperature and salinity contrast in the domain. Equations for the salinity and temperature contrasts in the fingering region were written down and combined with the equations for the motion of the salt fingers obtained by [Stern \(1960\)](#). The finger length was related to the vertical velocity by the continuity equation. These equations were combined to get an equation in terms of the growth rate. The maximum growth rate was then found, and this occurred at the wavelength of the fastest growing mode.

For the case of the domain used here, the height of the domain is 2 and the height of the fingering interface is 0.2. We also know that $R_T = 10^6$ and $R_\rho = 1.8$. Therefore the fastest growing wavelength is given by

$$\lambda = (2\pi) (2) \left[\left(\frac{2}{0.2} \right) (10^6) \left(1 - \frac{1}{1.8} \right) \right]^{-1/4} = 0.27. \quad (2.27)$$

In the previous sections a wavelength of 0.2 was used in order to reproduce [Özgökmen and Esenkov \(1998\)](#). This gave 5 waves in the 1 by 2 domain. [Özgökmen and Esenkov \(1998\)](#) states that wavenumbers of 5, 6 and 7 all yielded similar growth rates, so 5 as chosen.

2.5.5 Growing fingers with different wavelengths in ICOM

In the preceding simulations, the wavelength of the salt fingers comes from the initial condition. The initial condition was chosen because it was used by [Özgökmen and Esenkov \(1998\)](#), and the aim was to reproduce their results. The wavelength was chosen by [Özgökmen and Esenkov \(1998\)](#) because it is close to the fastest growing wavelength of the system. The growth rates for different wavelengths can be investigated by changing the wavelength of the initial sinusoidal perturbation. This should affect the growth rate of the salt fingers. It is also possible that the fastest growing mode could start to emerge even if the system was initialised with another wavelength.

Figure 2.12 shows the salinity fields for three different fixed runs where the initial conditions were changed so that the initial perturbations had different wavelengths: $\lambda = 0.1$, $\lambda = 0.2$ and $\lambda = 0.4$. An unstructured mesh with 150,000 nodes was used. All the modes exhibit the classic salt finger shape with plumes on the ends, with the plumes becoming more apparent as the simulation progresses.

Equation 2.26 gives 0.27 as the wavelength of the fastest growing mode. This is closest to the $\lambda = 0.2$ mode, so this one should grow the fastest according to theory. Figures 2.13(a) and 2.13(b) show that the $\lambda = 0.2$ mode is the fastest of the three as expected. The $\lambda = 0.1$ wave grows more slowly than the other ones because it is smaller than the fastest growing mode. The 0.4 mode is slower growing but in the simulation a faster growing mode overtakes it. The image for the $\lambda = 0.4$ mode (figure 2.12) shows that a smaller wave grows in the middle of each wave, and this grows faster than the $\lambda = 0.4$ wave. This smaller wavelength must be closer to the fastest growing mode.

From the standard deviation graph (figure 2.13(c)), it seems that the larger the initial wavelength the smaller the standard deviation, which correlates to a more developed secondary instability. Figure 2.12 clearly shows the $\lambda = 0.1$ and $\lambda = 0.2$ modes going unstable at $\lambda = 0.04$ s, when the $\lambda = 0.4$ mode is still stable. The finger shapes are still stable at 0.06 seconds, when the other two are very unstable. The fingers by this time have become thinner with large blobs in the ends. This mode has more space in




Figure 2.12: The salinity fields for three different initial conditions in a domain 4 units wide and 2 units high. The wavelengths of the initial perturbations were: $\lambda = 0.1$ (left), $\lambda = 0.2$ (middle) and $\lambda = 0.4$ (right). Time is increasing down the page for all three runs: 0, 0.01, 0.02, 0.03, 0.04, 0.05, 0.06. These are all fixed mesh runs with 150,000 nodes. The $\lambda = 0.2$ mode grows fastest as predicted by theory. The $\lambda = 0.4$ run shows another wavelength growing and taking over from the $\lambda = 0.2$ mode. This new wavelength has a similar growth rate to the $\lambda = 0.2$ mode. There is also a secondary instability that forms as time progresses: the fingers gradually become different lengths from their neighbours. This instability is much less apparent in the $\lambda = 0.4$ simulation because the fingers are further apart. This shows that the interaction between the fingers plays a crucial role in the formation of the secondary instability.

between the fingers, which means that it does not go unstable so rapidly because there is less interaction between the fingers.

Linear temperature and salinity gradients

The previous experiments have all started with a sharp interface. It is also possible to model salt fingers with gradients in temperature and salinity that span the whole height of the domain and without an initial condition to force the wavelength of the salt fingers so that the fastest growing mode emerges naturally. In this simulation the salinity and temperature were both initially 0 at the bottom and 1 at the top, with a linear gradient in between. There was a random number added to the initial condition at each point, so the fingers were not forced to grow with a particular wavelength. The random numbers were on a normal distribution around zero, with a standard deviation of 0.03. Since the salt fingers are not forced to grow in a particular mode, it should be possible to tell which is the fastest growing mode for the system by analysing the wavelength of the fastest growing disturbances in this run.

The salinity and temperature fields at 0.1s are shown in figure 2.14. The salt finger shape can clearly be seen in this image. There are now salt fingers throughout the domain, and there is no distinct fingering interface. The salinity and temperature anomalies are shown in figure 2.15. The initial condition was subtracted from the salinity and temperature fields so that the changes can be seen clearly. This shows that there has been fingering activity throughout the domain with regions of positive and negative anomalies at several points representing fingering activity.

Equation 2.26 gives a wavelength of 0.24 for the fastest growing mode in this case. This was calculated in a similar way to the previous case with the sharp interface, except the height of the fingering region is now 2, the height of the whole domain. A python function was used to find the contours and obtain a preliminary estimate for the fastest growing mode. At the contour salinity = 0.5 the fastest growing mode had a wavelength of about 0.27. This appears to be the right magnitude.



2.6 Discussion of results

Salt fingers have been modelled in two dimensions with using ICOM. The results for a periodic domain with a structured mesh are compared to the results of Özgökmen and Esenkov (1998) and ICOM is able to reproduce these results. They show close agreement in the growth rate and the flux ratio. An initial perturbation is then added to the salt fingers so that their initial lengths are different, the perturbation grows throughout the simulation. This shows that the salt fingers are unstable to small perturbations.


The instability is also present when an unstructured mesh is used. It was shown that the instability is greater for a lower resolution mesh because the differences between the length of the fingers was greater. A comparison between fixed and adaptive meshes was carried out using ICOM. The secondary instability was used as a measure of the effective resolution of the adaptive meshes. The adaptive mesh was shown to reproduce the same results as the fixed mesh using fewer nodes. Therefore the adaptive mesh technique leads to savings in computational cost without compromising the results.

The simulation was run with several different initial conditions with different wavelength for the initial displacement. There is a relationship between the wavelength of the initial disturbance and the growth rate. These results were in agreement with theoretical predictions of the fastest growing mode (Kunze, 1987).

2.6. DISCUSSION OF RESULTS



(a) The size of the upward and downward propagating fingers as a function of time for three different initial perturbations, with specified wavelengths. (b) The average velocities of the upward and downward propagating fingers against time.



(c) The deviations of the finger lengths as a function of time for three different initial perturbations, with specified wavelengths.

Figure 2.13: A comparison of salt finger growth for different initial perturbation wavelengths: $\lambda = 0.1$, $\lambda = 0.2$ and $\lambda = 0.4$. The $\lambda = 0.2$ run has the fastest growth rate as predicted by theory. The $\lambda = 0.4$ mode is slow growing, but a shorter mode grows out of the top of the waves, so the growth rate for this run is nearly as high as the growth rate for the $\lambda = 0.2$ run. See figure 2.12 for the salinity fields.




Figure 2.14: Salinity field for a fixed mesh run with 150,000 nodes with linear temperature and salinity gradients in a domain 4 wide and 2 high at time 0.1. The fingering activity started from a random perturbation.




Figure 2.15: Salinity and temperature anomalies for a fixed mesh run with 150,000 nodes with linear temperature and salinity gradients in a domain 4 cm wide and 2 cm high at 0.1s. The changes in salinity and temperature are easier to see in this graph. The fingering regions stand out clearly.

Chapter 3

The restratification of a water column following open ocean deep convection

3.1 Introduction

Open Ocean Deep Convection (OODC) is the convective mixing of dense water below the mixed layer. It occurs in particular regions of the North Atlantic: the Labrador, Greenland and Mediterranean seas as well as occasionally in the Weddell Sea. These regions have weak stratification and experience intense buoyancy loss to the atmosphere. Surface waters are mixed to a great depth via deep convection currents, forming convective chimneys. The convection sites in the north Atlantic, particularly the Labrador Sea help to feed the global Meridional Overturning Circulation ([Marshall and Schott, 1999](#)). This project will focus on the breakdown of the convection chimneys and restratification of the water in a region where deep convection has recently occurred.

The Meridional Overturning circulation (MOC) is a system of global ocean currents. These currents have a significant impact on our climate because they transport large quantities of heat around the globe. In particular, the Atlantic MOC transports warm water Northwards in the upper layers of the ocean and cooler water Southwards in the deep ocean. After being transported Northwards, the warm water sinks and cools in the deep convection regions in the North Atlantic, and is then transported South. Since the OODC is thought to help drive the MOC in the North Atlantic, a change to the OODC could affect the climate. If the ice caps melted or the hydrological cycle intensified as a result of global warming, there would be more fresh water on the surface of the ocean, which could slow or stop the OODC ([Challenor P.G and R., 2006](#)). This in turn could slow down the MOC which may have an impact on the climate. It has been suggested that the Northern Hemisphere could cool by an average of 1-2 °C and the southern hemisphere would warm slightly. There would also be a fall in rainfall in the Northern Hemisphere ([Vellinga and Wood, 2002](#)).

A weakening of the MOC has been observed already although it is not clear that this represents a long term trend, and there is a high degree of uncertainty in the results ([Quadfasel, 2005](#)). Model studies have shown that the MOC slows down in response to a fresh water flux. Stouffler ([Stouffer et al., 2006](#)) ran several models for 100 years with a 0.1 Sv ($1 \text{ Sv} = 10^6 \text{ m}^3 \text{ s}^{-1}$) fresh water flux added to a region of the North Atlantic. All the models show a slowing of the MOC but it is unclear how much it would slow down as there is substantial disagreement between models. None of the models predict a complete shut down of the MOC. Some of the models predict a re-intensification of the MOC when the convection sites move North to a region not affected by the hosing. This probably would not happen if the hosing region was bigger.




Figure 3.1: The phases of Open Ocean Deep Convection (image taken from [Marshall and Schott \(1999\)](#)): (a) preconditioning, (b) deep convection, (c) lateral exchange and spreading. Buoyancy flux through the sea surface is represented by curly arrows, the underlying stratification/outcrops is shown by continuous lines. The volume of fluid mixed by convection is shaded.

3.2 The process of Open Ocean Deep Convection

3.2.1 The phases of open ocean deep convection

Figure 3.1 shows the three main phases of open ocean deep convection. The first phase, preconditioning, occurs over a large area of ocean, of order 100 km. The gyre-scale cyclonic circulation causes “doming” of the isopycnals in the centre of the gyre. The denser mixed water is displaced upwards, nearer to the surface. (figure 3.1 (a)). Due to the particular meteorology of the area (eg surface cooling from wind), buoyancy forcing then triggers convection ([Marshall and Schott, 1999](#)).

The second phase is known as the convection or violent mixing phase (figure 3.1 (b)). During winter, the buoyancy loss erodes the near-surface stratification over a region of several hundred kilometres, so that the weakly stratified water is exposed to the surface forcing. Then subsequent cooling initiates convection in which a substantial part of the fluid column overturns in various plumes, particularly in the centre of the preconditioned dome, distributing the dense water vertically. Each plume is about 1 km or less in both the vertical and horizontal directions. The vertical velocities reach 10 cm s^{-1} ([Marshall and Schott, 1999](#)). The plumes together mix the properties of the whole area, forming a dense well mixed patch with a diameter which can range from tens of kilometres to over 100 km. The mixed patch can form in two different ways. In the first, known as a chimney, the plumes cause a net downward transport of dense water. In the second, the downward motion of the plumes is compensated by upward motion between the plumes, so that the resultant mean vertical velocity is zero. In this case the dense patch is formed by mixing of the cold surface waters with warmer waters below ([Lavender and Davis, 2002](#); [Marshall and Schott, 1999](#); [Killworth, 1982](#)).

The final phase phase is known as the restratification or lateral exchange phase (figure 3.1 (c)) ([Marshall and Schott, 1999](#)). The vertical heat transfer on convective scales ceases when the forcing stops, or after many days of cooling. Horizontal heat transfer then occurs on on geostrophic scales. The mixed patch exchanges fluid with the surrounding water in an eddying motion. Over a period of a few months, the mixed fluid gradually restratifies with the surroundings due to gravity and rotation.

3.2.2 The Labrador Sea

The deep convection regions in the North Atlantic are particularly important because they play a role in maintaining the Meridional Overturning Circulation. The Labrador sea is one of these. During the preconditioning phase there is a cyclonic motion which forms the doming isopycnals, as mentioned

above. At this stage, there are three main layers. The top layer is fresh and formed as a result of eddy transport from the boundary currents at the edge of the Labrador Sea. At 200 – 700 m there is a weak temperature and salinity maximum due to the Irminger Sea Water. From 700 – 2000 m there is a layer of Labrador Sea Water, left from the previous winter’s convection. This recirculates in the western basin. At the bottom is cold, dense Denmark Strait Overflow Water. This also circulates cyclonically. There is another intermediate salinity maximum at 2500 – 3200 m. This water is from the Gibbs Fracture Zone and Eastern Basin Mediterranean Water (Marshall and Schott, 1999). The preconditioning phase was observed in the Labrador sea by Lavender and Davis (2002). They found direct evidence for a cyclonic gyre in a region of weak flow in the western basin. Because of the weak flow, the water remains in the region for a long time, and is subjected to a large wintertime heat loss which eventually triggers the deep convection phase.

Deep convection has been inferred by several measurements, although the lateral extent of the convection is not known (Marshall and Schott, 1999). The presence of deep convection in the Labrador Sea was inferred from hydrographic measurements showing that the well-mixed layer penetrates to depths of 2300 m (Dickson et. al. 1996). The measurements also found a region in the Western Labrador sea where surface waters are denser than the surroundings (Lazier 1973). The deep convection has rarely been observed directly. Gascard and Clarke (1983/1983) and Lilly et al (1999) both observed vertical velocities of up to 9 cm s⁻¹ in winter, when these were strong winds and the air temperature was cold. These studies were limited because they only took place over small areas of ocean (Lavender and Davis, 2002). Gascard and Clarke found a 50 km wide area of homogeneous Labrador deep water. They also observed horizontal eddying on several different length scales of 5 to 50 km, with the convection itself occurring on length scales of 1 km (Killworth, 1982). In the winters of 1996/97 and 1997/98, the Labrador Sea Deep Convection Experiment (Lab Sea Group 1998) conducted observations and produced models of OODC. The results from the observation part of the experiment are discussed in Lavender and Davis, 2002 (Lavender and Davis, 2002).

3.3 Modelling the restratification of a convection chimney

Rousset (Rousset et al., 2009) carried out a multi-model study of the restratification phase of a convection chimney. They observed the restratification process, including the eddies that form due to baroclinic instability. This process has also been modelled by Jones and Marshall (1997), and eddies were also observed in this case. The three models compared by (Rousset et al., 2009) were an isopycnic model with a mixed layer (the Miami Isopycnic Coordinate Ocean Model; MICOM); a version of MICOM with a deep adiabatic layer outcropping at the surface instead of a mixed layer (MICOM-ADIAB) and OPA, a z-coordinate model. All three models solve the three-dimensional hydrostatic primitive equations on a staggered horizontal C-grid with no-slip conditions at the solid boundaries. The salinity is constant in all the models.

3.3.1 The Three Models

The OPA model is the only one that has the vertical coordinate in z , with a free-surface on top. Temperature is calculated by solving the advection-diffusion equation with a second order centred-difference scheme. The hydrostatic instability is treated with non-penetrative convective adjustment scheme. The viscosity and diffusion are parameterised by second order diffusive operators with constant diffusion coefficients.

MICOM has density as the vertical coordinate. There is a vertically homogeneous, density varying mixed layer which deals with diabatic surface exchanges with the atmosphere. Vertical mixing with the mixed layer is modelled by an entrainment equation (Kraus and Turner, 1967). Under the mixed layer, the ocean is modelled with discrete layers of constant density. The interaction between the layers occurs via hydrostatically transmitted pressure torques. The stratification changes when individual layer thicknesses change, as described by the continuity (mass conservation) equation. There is no mass-transfer allowed across the isopycnals, except between the mixed layer and the underlying ocean. This happens when hydrostatic instabilities occur in the mixed layer as a result of the properties of the mixed layer changing. Inside the mixed layer, advection and diffusion of temperature and momentum are only allowed to occur horizontally. The second order diffusive operators parameterise mixing which occurs on

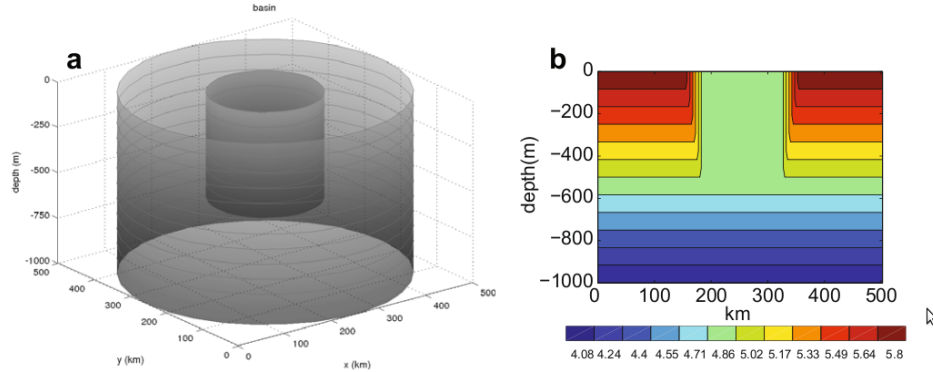


Figure 3.2: The initial configuration of the experiments discussed in [Rousset et al. \(2009\)](#) (figure taken from [Rousset et al. \(2009\)](#)). The model domain (a) is a cylinder of radius $R = 250$ km and depth $H_b = 1$ km. The initial stratification (b) is 12 equally spaced layers which extend over the whole of the domain except for a central region of radius $R_c = 70$ km and depth $H_c = 580$ m, which initially has a constant density.

smaller scales than the grid spacing. The scalar properties all have a constant diffusivity but momentum has a deformation dependent viscosity ([Smagorinsky, 1963](#)). MICOM-ADIAB's mixed layer is a layer of constant density the same as all the others. Each layer conserves mass except for a weak prescribed diapycnal flux (flux normal to the isopycnals).

3.3.2 Experimental set-up

Because salinity is constant, density ρ is dependent only on temperature T . In the equation of state, the density is linearised around an initial around the initial bottom temperature using α . The initial configuration is representative of the conditions at the end of a convective period in a sub-polar convective basin. The stratification is typical of the Labrador sea. The whole domain is a cylinder. Initially there is a cylinder in the middle with constant temperature and a cylinder around the edges with twelve layers of constant temperature and a constant buoyancy frequency. The region between the two cylinders is linearly stratified in temperature. The area of constant density extends over the top seven layers in OPA, and is contained in the mixed layer of the other models. There is no mechanical or buoyancy source of turbulent energy at the surface so mixed layer entrainment or retreat cannot occur. The stratification in the outer region is representative of the stratification in the Labrador sea, excluding the Polar Surface Water, which is very light, and the linearly varying rim is also the same as the observed conditions. However, the temperature is not the same as the observed temperature profile because the salinity is constant in the simulation. The parameters are shown in table 3.1. The horizontal resolution is 2.5 km. The diffusion timescale across the basin is large compared to the restratification timescale. The isopycnal models may have locally larger viscosity.

3.4 Future Plans

The restratification phase of open ocean deep convection will be modelled with particular emphasis on the vortices that occur due to the baroclinic instability. The aim is to model one of these vortices interacting with a wall or with each other. A detailed study of this will be carried out. ICOM should provide an advantage over other models because the adaptivity will allow the vortices to be modelled in higher resolution. This will be quantified by comparing the fixed and adaptive meshes.


The first step is to investigate the structure of the vortices by reproducing the results of [Rousset et al. \(2009\)](#), described in the chapter 3.3, using ICOM. The problem will be set in the way described in the paper and the results from ICOM will be compared to those in [Rousset et al. \(2009\)](#). Initially a fixed mesh will be used, and when the technique has been perfected an adaptive mesh will be used. The adaptive mesh should improve the efficiency, and the vortices observed by [Rousset et al. \(2009\)](#) should be resolved in more detail. This leads into the detailed study of the vortices.

3.4. FUTURE PLANS

Table 3.1: A summary of the parameters used in the simulations analysed by Rousset ([Rousset et al., 2009](#))

Parameter	Symbol	Value
Radius of the domain	R	250 m
Vertical height of the domain	H_b	1 km
Coriolis Parameter	f_0	10^{-4} s^{-1}
Salinity	S	34.9 pss (held constant)
Thermal expansion coefficient	α	$1.95 \times 10^{-4} \text{ }^\circ\text{C}^{-1}$
Radius of central homogeneous cylinder	R_c	70 km
Height of central homogeneous cylinder	H_c	580 km (to surface)
Temperature of central homogeneous cylinder	T_c	4.86°C
Buoyancy frequency of outer layer	N	$1.9 \times 10^{-3} \text{ s}^{-1}$
Thickness of each layer of constant temperature in the outside region	h_0	83.33 m
Horizontal thermal diffusivity, viscosity	κ_T, ν	$7.5\text{m}^2\text{s}^{-1}$

The Gantt chart (figure 3.4) shows the tasks to be completed in order to finish the PhD. The time is measured in months from February 2010.



Bibliography

- H. K. Challenor P.G and M. R. *Towards the Probability of Rapid Climate Change*, chapter 7, pages 55–63. Cambridge University Press, 2006.
- I. Dmitrenko, S. Kirillov, V. Ivanov, and R. Woodgate. Mesoscale Atlantic water eddy off the Laptev Sea continental slope carries the signature of upstream interaction. *Journal of Geophysical Research (Oceans)*, 113(12):7005, July 2008.
- R. Ford, C. Pain, M. Piggott, A. Goddard, C. de Oliveira, and A. Umpleby. A nonhydrostatic finite-element model for three-dimensional stratified oceanic flows. Part I: model formulation. *Monthly Weather Review*, 132(12):2816–2831, 2004.
- M. Gregg. The study of mixing in the ocean: A brief history. *Oceanography*, 4(1):39–45, 1991.
- J. Holyer. On the collective instability of salt fingers. *Journal of Fluid Mechanics Digital Archive*, 110: 195–207, 1981.
- H. Jones and J. Marshall. Restratification after deep convection. *Journal of Physical Oceanography*, 27(10):2276–2287, 1997.
- P. Killworth. Deep convection in the world ocean. *Reviews of Geophysics*, 21(1), 1982.
- E. Kraus and J. Turner. A one-dimensional model of the seasonal thermocline: II. The general theory and its consequences. *Tellus*, 19(1):98–105, 1967.
- E. Kunze. Limits on growing, finite-length salt fingers: A Richardson number constraint. *Journal of Marine Research*, 45(3):533–556, 1987.
- K. Lavender and R. Davis. Observations of Open-ocean Deep convection in the Labrador Sea from subsurface floats. *Journal of Physical Oceanography*, 32(2):511–526, 2002.
- J. Ledwell and A. Watson. Evidence for slow mixing across the pycnocline from an open-ocean tracer-release experiment. *Nature*, 364(6439):701–703, 1993.
- J. Marshall and F. Schott. Open-ocean convection: Observations, theory, and models. *Rev. Geophys*, 37(1):1–64, 1999.
- C. J. E. Morell, J. M. and W. J. Merryfield. Thermohaline staircases in a Caribbean eddy and mechanisms for staircase formation. *Deep Sea Research Part II: Topical Studies in Oceanography*, 53:128–139, Jan. 2006. doi: 10.1016/j.dsr2.2005.09.013.
- W. Munk. Abyssal recipes. *Deep-Sea Res*, 13(4), 1966.
- T. Özgökmen and O. Esenkov. Asymmetric salt fingers induced by a nonlinear equation of state. *Physics of Fluids*, 10(8):1882–1890, 1998. doi: 10.1063/1.869705. URL <http://link.aip.org/link/?PHF/10/1882/1>.
- C. C. Pain, A. P. Umpleby, C. R. E. de Oliveira, and A. J. H. Goddard. Tetrahedral mesh optimisation and adaptivity for steady-state and transient finite element calculations. *Computer Methods in Applied Mechanics and Engineering*, 190(29-30):3771 – 379, 2001.

BIBLIOGRAPHY

- S. A. T. J. Piasek. Nonlinear evolution and structure of Saltfingers, Elsevier Oceanography Series. *Marine Turbulence*, 28:193–219, 1980.
- D. Quadfasel. The Atlantic heat conveyor slows. *Nature*, 438(1):565–566, 2005.
- C. Rousset, M. Houssais, and E. Chassignet. A multi-model study of the restratification phase in an idealized convection basin. *Ocean Modelling*, 26(3-4):115–133, 2009.
- R. Schmitt. The growth rate of super-critical salt fingers. *Deep Sea Research Part I: Oceanographic Research*, 26:23–40, Jan. 1979. doi: 10.1016/0198-0149(79)90083-9.
- R. W. Schmitt, J. R. Ledwell, E. T. Montgomery, K. L. Polzin, and J. M. Toole. Enhanced Diapycnal Mixing by Salt Fingers in the Thermocline of the Tropical Atlantic. *Science*, 308(5722):685–688, 2005.
- C. Shen and G. Veronis. Numerical simulation of two-dimensional salt fingers. *Journal of Geophysical Research*, 102:23,13123,143, 1997.
- J. Smagorinsky. General circulation experiments with the primitive equations. *Monthly weather review*, 91(3):99–164, 1963.
- M. Stern. The salt fountain and thermohaline convection. *Tellus*, 12(2):172–175, 1960.
- R. Stewart. *Introduction to Physical Oceanography*. 2005.
- H. Stommel. An oceanographical curiosity: the perpetual salt fountain. *Deep Sea Research*, 3(2):152–153, 1953.
- R. Stouffer, J. Yin, J. Gregory, K. Dixon, M. Spelman, W. Hurlin, A. Weaver, M. Eby, G. Flato, H. Hasumi, et al. Investigating the causes of the response of the thermohaline circulation to past and future climate changes. *Journal of Climate*, 19(8):1365–1387, 2006.
- J. Taylor and G. Veronis. Experiments on double-diffusive sugarsalt fingers at high stability ratio. *Journal of Fluid Mechanics*, 321:315–333, 1996.
- M. Vellinga and R. Wood. Global Impacts of the Collapse of the Atlantic Thermohaline Circulation. *Climatic Change*, 54:251–267, 2002.
- D. Whitfield, G. Holloway, and J. Holyer. Spectral transform simulations of finite amplitude double-diffusive instabilities in two dimensions. *Journal of Marine Research*, 47(2):241–265, 1989.
- J. Yoshida and H. Nagashima. Numerical experiments on salt-finger convection. *Progress in Oceanography*, 56:435–459, 2003.

Chapter 4

Appendix

4.1 Setting up the mesh

The mesh was produced using Gmsh, producing a .geo file which describes the geometry of the domain: it contains the coordinates of the box and the spacing of the nodes and identifies physical groups (edges or surfaces) that need to be referenced in the code, e.g. to apply boundary conditions; and a .msh file which contains the output mesh. This is then converted to triangle format for input to fluidity, which gives three files: .node, .ele and .edge (a .face is used instead of .edge in 3D)

4.2 Non-dimensional parameters

The governing parameters for salt fingering are:

The Prandtl number,

$$Pr = \nu / K_T \quad (4.1)$$

The Lewis number

$$\tau = \frac{K_S}{K_T} \quad (4.2)$$

The thermal Rayleigh number

$$R_T = \frac{g\alpha\Delta TL^3}{K_T\nu} \quad (4.3)$$

The saline Rayleigh number

$$R_S = \frac{g\beta\Delta SL^3}{K_T\nu} \quad (4.4)$$

Schmitt number

$$Sc = \nu / K_S \quad (4.5)$$

The density ratio

$$R_\rho = \frac{\alpha\Delta T}{\beta\Delta S}, \quad (4.6)$$

The Stern number

$$St = \frac{\beta F_S - \alpha F_T}{\nu(\alpha T_z - \beta S_z)} \quad (4.7)$$

salt Nusselt number

$$N_S = \frac{F_S\Delta S}{\kappa_S L} \quad (4.8)$$

thermal Nusselt number

$$N_T = \frac{F_T\Delta T}{\kappa_T L} \quad (4.9)$$

where ν is the molecular diffusivity of momentum, F_S and F_T are the diapycnal salt and heat fluxes, T_z and S_z are the heat and salt gradients, L is the salt finger layer depth - normally the height of the

4.2. NON-DIMENSIONAL PARAMETERS

domain in a single layer experiment or the layer height if there are multiple layers. α is the linear thermal expansion coefficient and β is the linear saline contraction coefficient, given by

$$\alpha = -\frac{1}{\rho} \frac{\partial \rho}{\partial T} \quad \beta = \frac{1}{\rho} \frac{\partial \rho}{\partial S}, \quad (4.10)$$